

AMENDMENTS TO THE CLAIMS

1. (Cancelled)
2. (Currently Amended) A method of delivering content across a plurality of zones within a network, comprising:
 - receiving a request from a client located within one of the plurality of zones for access to resources associated with a domain name;
 - determining network conditions for the network based on a determination of the load for each of the plurality of zones;
 - distributing the request to one of the plurality of zones based on the determined network conditions;
 - selecting one of a plurality of servers within the zone in which the request was distributed, the selection of the server being based on a determination for optimally balancing the load across the plurality of servers;
 - resolving an Internet protocol (IP) address of the selected server; and
 - determining whether to delegate delivery of the resources to a content delivery network based on the determination for optimally balancing the load across the plurality of servers, a Time to Live value set at a pool level for each of a plurality of available pools, and a pool load-balancing setting.
3. (Previously presented) The method of Claim 2, further comprising querying a local Domain Name System (DNS) to provide the IP address associated with the domain name.
4. (Previously presented) The method of Claim 3 wherein when the IP address is not present at the local DNS, querying a primary DNS to resolve the IP address associated with the domain name.
5. (Previously presented) The method of Claim 4, wherein when the primary DNS determines the domain name is delegated to a EDNS, further comprises referring the local

DNS to the EDNS to resolve the IP address for the selected server, the EDNS employs at least one of a plurality of load balancing determinations to select one of the plurality of servers and resolve the IP address for the selected server.

6–7. (Canceled)

8. (Previously presented) The method of Claim 2 wherein selecting one of the plurality of servers, further comprises:

- marking each of a plurality of pools to a not tried state;
- determining the pool load-balancing setting;
- selecting one of the plurality of pools that is marked to the initialization state;
- marking the selected one of the plurality of pools to a tried state;
- attempting to obtain an answer using the determined pool load-balancing on the selected one of the plurality of pools;
- determining if the answer was obtained.

9. (Previously presented) The method of Claim 5 wherein selecting one of the plurality of servers, further comprises accounting for computing power behind a wide area virtual server.

10–16. (Canceled)

17. (Previously presented) The method of Claim 2, further comprising:
deriving cost metrics for network paths in topological maps;
using the cost metrics to determine a geographic location of the request; and
distributing the request based on the geographic location.

18. (Previously presented) The method of Claim 2, further comprising:
using a last hop address to determine a geographic location of the request; and
distributing the request based on the geographic location.

19. (Previously presented) The method of Claim 2, further comprising using a last-resort pool to select the server.

20. (Previously presented) The method of Claim 2, wherein at least a portion of the plurality of servers are virtual servers.

21. (Previously presented) The method of Claim 20 wherein selecting one of the plurality of servers within the zone in which the request was distributed, the selection of the server being based on a determination for optimally balancing the load across the plurality of servers, further comprising:

determining if the selected server is a virtual server, and if so:

determining a number of nodes up on the virtual server;

determining if the number of nodes up or the number of connections for the virtual server exceeds a predetermined number and if so returning a value indicating the capacity of the virtual server has been exceeded.

22. (Currently Amended) A system for balancing the load on a plurality of virtual servers that provide access to resources associated with a domain name, comprising:

a memory for storing logical instructions; and

a processor for executing the logical instructions stored in the memory, the execution of the logical instructions causing functions to be performed, including:

receiving a request from a client located within one of the plurality of zones for access to resources associated with a domain name through an authoritative server;

determining network conditions for the network based on a determination of the load for each of the plurality of zones;

distributing the request to one of the plurality of zones based on the determined network conditions;

selecting one of the plurality of servers within the zone in which the request was distributed, the selection of the server being based on a determination for optimally balancing the load across the plurality of servers;

resolving an Internet protocol (IP) address of the selected server; and determining whether to delegate delivery of the resources to a content delivery network based on the determination for optimally balancing the load, wherein the determination of whether to delegate delivery is made prior to a connection of the client to the content delivery network if the delivery of resources is delegated to a delegation pool associated with the content delivery network.

23. (Previously presented) The system of Claim 22, wherein selecting one of the plurality of servers, further comprises choosing the server based on one of a plurality of static load balancing determinations for each server, the plurality of static load balancing determinations being selectable and including random, round robin, static ratio, global availability and topology.

24. (Previously presented) The system of Claim 22 wherein selecting one of the plurality of servers, further comprises choosing the server based on one of a plurality of dynamic load balancing determinations for each server, the dynamic load balancing determinations being selectable and including completion rate, least connections, packet rate, hops, round trip times, new connection rate, kilobyte rate, quality of service and dynamic ratio.

25. (Previously presented) The system of Claim 22 wherein selecting one of the plurality of servers, further comprises:

marking each of a plurality of pools to a not tried state;
determining a pool load-balancing setting;
selecting one of the plurality of pools that is marked to the initialization state;
marking the selected one of the plurality of pools to a tried state;
attempting to obtain an answer using the determined pool load-balancing on the selected one of the plurality of pools;

determining if the answer was obtained.

26. (Previously presented) The system of Claim 22 wherein selecting one of the plurality of servers, further comprises accounting for computing power behind a wide area virtual server.

27-28. (Cancelled)

29. (Previously presented) The system of Claim 22, further comprising distributing the request to a closest server based on a closest identified geographic CONTENT-SERVER or a best-performing CONTENT-SERVER, or a best quality path.

30-32. (Cancelled)

33. (Previously presented) The system of Claim 22, further comprising: deriving cost metrics for network paths in topological maps; using the cost metrics to determine a geographic location of the request; and distributing the request based on the geographic location.

34. (Previously presented) The system of Claim 22, further comprising: using a last hop address to determine a geographic location of the request; and distributing the request based on the geographic location.

35. (Previously presented) The system of Claim 22, further comprising using a last-resort pool to select the server.

36. (Previously presented) The system of Claim 22 wherein at least a portion of the plurality of servers are virtual servers.

37. (Previously presented) The system of Claim 36 wherein selecting one of the plurality of servers within the zone in which the request was distributed, the selection of the server being based on a determination for optimally balancing the load across the plurality of servers further comprising:

determining if the selected server is a virtual server, and if so:

determining a number of nodes up on the virtual server;

determining if the number of nodes up or the number of connections for the virtual server exceeds a predetermined number and if so returning a value indicating the capacity of the virtual server has been exceeded.

38. (Currently Amended) A machine-readable storage medium having machine-executable instructions stored thereon, which when executed by at least one processor, causes the at least one processor to perform one or more actions for enabling delivery of content across a plurality of zones within a network, the actions comprising:

receiving a request from a client located within one of the plurality of zones for access to resources associated with a domain name;

determining a physical geographic location associated with the internet protocol (IP) address of the request;

determining network conditions for the network;

distributing the request to one of the plurality of zones based on the geographic location and network conditions;

selecting one of the plurality of servers within the zone in which the request was distributed, the selection of the server being based on a determination for optimally balancing the load across the plurality of servers;

resolving an Internet protocol (IP) address of the selected server; and

determining whether to delegate delivery of the resources to a content delivery network based on the determination for optimally balancing the load.

39. (Previously presented) The machine-readable storage medium of Claim 38 wherein selecting one of the plurality of servers further comprises choosing the server based on one of a plurality of static load balancing determinations for each server, the plurality of static load balancing determinations being selectable and including random, round robin, static ratio, global availability and topology.

40. (Previously presented) The machine-readable storage medium of Claim 38 wherein selecting one of the plurality of servers further comprises choosing the server based on one of a plurality of dynamic load balancing determinations for each server, the dynamic load balancing determinations being selectable and including completion rate, least connections, packet rate, hops, round trip times, new connection rate, kilobyte rate, quality of service and dynamic ratio.

41. (Previously presented) The machine-readable storage medium of Claim 38, further comprising:

determining a geographic location of the request;
identifying a proximity of the request based on the geographic location; and
distributing the request based on the proximity.

42–44. (Canceled)

45. (Previously presented) The machine-readable storage medium of Claim 41 wherein determining the geographic location of the request further comprises using a last hop address to determine the geographic location of the request.

46. (Previously presented) The machine-readable storage medium of Claim 38, further comprising using a last-resort pool.

47. (Previously presented) The machine-readable storage medium of Claim 38 wherein at least a portion of the plurality of servers are virtual servers.

48. (Previously presented) The machine-readable storage medium of Claim 47 wherein selecting one of the plurality of servers within the zone in which the request was distributed, the selection of the server being based on a determination for optimally balancing the load across the plurality of servers further comprising:

determining if the selected server is a virtual server, and if so:

determining a number of nodes up on the virtual server;

determining if the number of nodes up or the number of connections for the virtual server exceeds a predetermined number and if so returning a value indicating the capacity of the virtual server has been exceeded.

49. (Currently Amended) A system for delivering content across a plurality of zones within a network, comprising:

a receiver that receives a request from a client located within one of the plurality of zones to access resources associated with a domain name;

a means for distributing the request to one of the plurality of zones based on network conditions of the plurality of zones;

a load balancer that selects one of a plurality of servers within the zone in which the request is distributed;

a system component that resolves an Internet protocol (IP) address of the selected server; and

a means for determining whether to delegate delivery of the resources to a content delivery network based on a load balancing determination for optimally balancing the load across the plurality of servers or the content delivery network, wherein the plurality of servers and the content delivery network are designated members of different delegation pools of content sources.

50. (Currently Amended) A method for delivering content over a network, comprising:

receiving a request from a client for access to resources associated with a domain name;

distributing the request to one of a plurality of zones based at least in part on network conditions associated with the zones, wherein distributing the request includes determining whether to delegate delivery of the resources to a content delivery network;

selecting one of a plurality of servers within the zone in which the request is distributed, wherein selecting the server is based at least in part on a balancing of the load among the servers;

resolving the domain name into an Internet protocol (IP) address of the selected server; and

determining whether to delegate, in a controlled manner, delivery of the resources to a content delivery network or to an origin site based on the balancing of the load and a mapping created between pools of content sources and an identified physical geographic location of an IP address of the request.